Tactical Immunized Maneuvering System for Exploration Air Vehicles

John Kaneshige* and K. Krishnakumar†
NASA Ames Research Center, Moffett Field, CA, 94035

Traditional flight management systems are designed to perform the necessary flight planning, navigation and guidance functions for flying from an origin to a destination. Special pre-constructed procedures and guidance modes are used for departing and arriving at airports, and performing special maneuvers such as holding patterns. However during an exploration mission, an aircraft must be capable of responding to feedback from the environment. As a result, it may be necessary to construct special procedures "on the fly", which take into account the dynamics and performance limitations of the aircraft. This paper describes a tactical maneuvering system that uses an artificial immune system based approach for constructing maneuver sequences. This approach incorporates the problem solving abilities and memory retention characteristics of an immune system. The resulting system is capable of making time-critical decisions in complex situations to accomplish near-term objectives within a dynamic environment. Simulation results demonstrate the potential of using immunized sequence selection in support of exploration missions using aerial vehicles.

I. Introduction

UNMANNED Aerial Vehicles (UAVs) have demonstrated potential as being effective platforms for supporting scientific and exploratory missions.[1] They are capable of performing long endurance flights, and reaching remote areas that may be too dangerous for humans. As their role and types of missions expand, from requiring remotely controlled to semi-autonomous and autonomous operations, challenges are presented which require onboard systems to have increasingly higher levels of intelligence.[2] These intelligent systems must be capable of making reliable decisions under varying conditions. As a result, they must incorporate the experience, reasoning and learning abilities of a pilot.

In terms of achieving a flight-path goal, a pilot's behavior can be captured through a layered model consisting of discrete-time strategic planning and tactical maneuvering, and continuous-time manual control (Fig. 1).[3] The discrete nature of strategic and tactical behaviors allows for automated decision-making techniques to be applied. Furthermore, since strategic planning decisions are less time-critical, more computationally intensive approaches can be utilized. All of the continuous-time processing elements can be isolated in the automation of manual control.

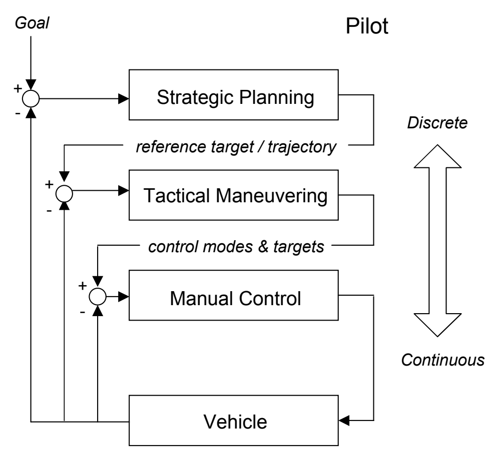

Figure 1. Pilot Behavior Hierarchy.

One aspect of exploration missions that differs from traditional flight operations is that the primary goal is to collect data, verses flying from one point to another. In many cases, the vehicle will need to react to the data that is being collected, instead of flying a pre-determined flight plan. Mission-specific "payload" sensors and data requirements can also result in different constraints being placed on how the vehicle must fly while collecting data on "targets" of interest. In data rich environments, there will be a number of tradeoffs that need to be taken into consideration in order to maximize both the quantity and quality of the data collected. These tradeoffs will include how quickly to "service" a target and at what level of resolution.

* Computer Engineer, Intelligent Systems Division, Mail Stop: 269-1, Member AIAA.
† Research Scientist, Intelligent Systems Division, Mail Stop: 269-1, Associate Fellow AIAA.

This paper provides an overview of how intelligent systems can be used to enhance exploration missions, and identifies some of the challenges that must be addressed. Particular emphasis is placed on the role of tactical maneuvering and how it can be used in conjunction with other intelligent systems to address many of the challenges. This paper also contains an overview of an intelligent tactical maneuvering system, along with an implementation description and simulation test results.

II. Enhancing Exploration Missions

A traditional Flight Management System (FMS) is designed to fly from an origin to a destination. Pilots or ground-based operators can enter a flight plan consisting of "waypoints" to establish an earth-based trajectory. Special lateral and vertical path tracking control laws are used to fly the "legs" between these waypoints. When it is necessary to deviate from the active flight plan, a new flight plan can be uploaded or commands can be sent to the automatic pilot (or autopilot) to fly a heading, altitude, airspeed, etc.

In the case of UAV exploration missions, the primary objective is to collect data rather than fly from one point to another. As a result, the mission can be enhanced if scientists and/or ground-based operators can redirect the flight based on feedback from payload sensors. However, one of the major problems is that communication latencies and human response times will often prohibit the timely response that is necessary for payload directed flight. As a result, intelligence must be incorporated into the onboard systems in order to "close-the-loop" around the payload sensors.

A. Roles of Onboard Intelligence

Strategic Planning:

In terms of the Pilot Behavioral Hierarchy (Fig. 1), when ground-based operators upload new flight plans they are essentially maintaining the role of strategic planner. However this responsibility can be passed on to onboard strategic planning systems that are capable of computing their own flight plans. Some of these flight planners have demonstrated the ability to re-plan in the presence of obstacles, using techniques such as evolutionary algorithms[4] and Voronoi diagrams[5]. These types of planners can also be used in support of exploration missions by re-planning to fly towards targets of interest, once the payload sensors have detected them.

Tactical Maneuvering:

The role of tactical maneuvering is normally performed, in a limited fashion, by simple heuristics within the FMS, which determine when it is necessary to switch from one control mode to another. However, when flying a Standard Instrument Departure (SID) or Standard Terminal Arrival Route (STAR), scripted procedures are used to transition between motion-based legs, such as flying a heading to an altitude. These procedures are designed to expedite Air Traffic Control (ATC) clearance delivery and to facilitate the transition to and from "enroute" operations. However, they also serve as a demonstration of how motion-based plans can be used to generate aircraft-centric trajectories, which are based on the maneuvering capabilities of the vehicle.

During tactical maneuvering, pilots use their knowledge of aircraft capabilities and near-optimal maneuvering strategies in order to select the necessary actions. These actions can be approximated by piece-wise linear or piece-wise constant commands, and switching between commands.[6] The interconnection of a finite number of commands can be used to generate motion-based plans that can exploit the full maneuvering capabilities of the aircraft.[7] Figure 2 shows an example of an aerobatic maneuver that can be performed using a motion-based plan. This plan consists of a sequence of aircraft-centric commands such as: *increase thrust, maintain zero roll rate, control pitch rate (to desired normal acceleration), capture –45 degree flight-path angle, bank (from –180 degrees) to zero degrees, and then capture desired altitude and airspeed.*

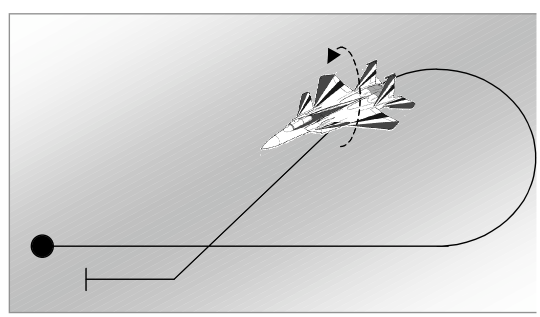

Figure 2. Half Cuban Eight Maneuver.

In terms of supporting exploration missions, motion-based plans are necessary when constraints have been placed on how the vehicle must fly while servicing a target. These constraints can depend on the capabilities of the payload sensors as well as the mission-dependent data requirements. For example, in some cases the vehicle may need to fly wings-level over a ground-based target, even under crosswind conditions. If sun angle needs to be taken into account, the vehicle may also need to be heading in a particular direction. In other cases, the vehicle may need to "point" the payload sensors at a target for a period of time (Fig. 3). In these situations the attitude of the vehicle may need to be constrained so that the target stays in the view of the payload sensors. In the event of a moving ground-based or airborne target, the relative motion of the target also needs to be taken into account. In some cases, the payload sensors may only return partial information, requiring the vehicle to fly in the general direction of a target without knowing its precise location. Once the location is realized the vehicle may be very close to the target, requiring semi-aggressive corrections to be executed in a time-critical fashion.

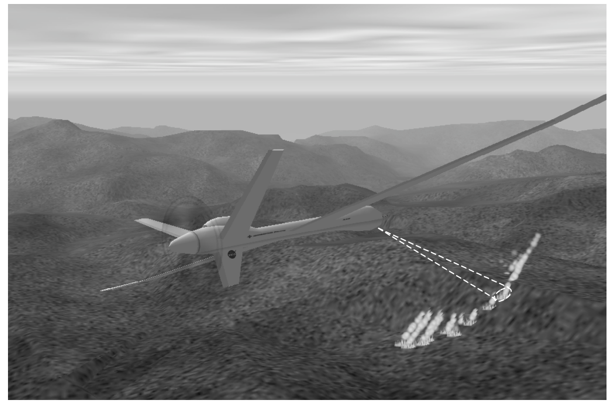

Figure 3. Target Pointing.

Automated Control:

The FMS and autopilot control laws are used to automate the manual control task. When the FMS is engaged, the vehicle is controlled along the specified fixed-based trajectory. When an autopilot control mode and target is specified, such as *bank* 10^o, the vehicle is controlled to the commanded state, resulting in a motion-based path. One method of further supporting exploration missions is to develop custom control laws to directly close the loop around the various payload sensors. These control laws would allow the vehicle to fly along the path of greatest data return, assuming that path is both identifiable and flyable. While this method requires continuous data feedback with minimum latency, it allows the vehicle to directly react to the data as it is being collected.

B. Technology Challenges

In order to maximize both the quantity and the quality of the data collected, there will be a number of tradeoffs that need to be taken into consideration. For example, when servicing a target it may be desirable to fly low enough to achieve a certain data resolution. However this may mean that the vehicle cannot capture the entire target with a single pass. As a result, the vehicle may need to make multiple passes, resulting in additional time required for servicing the target. In a target rich environment it may also be possible to service multiple targets at the same time or in direct succession. For all of these cases, one of the primary challenges is to come up with cost and benefit metrics in order to evaluate the success of a particular mission. Another major challenge is to be able to generate the plans that are necessary for maximizing the success of the mission.

These types of tradeoffs can occur at both the strategic and tactical planning levels. The primary difference is the mission scope and the timeline. While strategic planning focuses on the success of a mission as a whole, tactical maneuvering focuses more on the near-term objectives of a segment of the mission. As a result, tactical maneuvering "problems" will typically have fewer targets to consider, but will tend to require higher resolution trajectory "solutions". Furthermore these solutions will usually have to be computed in a timely fashion, since the vehicle may need to make quick corrections to a target once it becomes "visible" to the payload sensors.

A prohibitive factor in developing maneuvering systems in the past have stemmed from the varying nature of the problems that are encountered. On one hand numerous simple problems will often be encountered, which could be solved very quickly by a rule-based system. On the other hand very complex problems may also be encountered, which require more sophisticated search methods. This paper investigates the use of an Artificial Immune System (AIS) based approach for selecting maneuver sequences.[8] This approach takes advantage of the memory retention and adaptability characteristics of the biological immune system. The resulting system is capable of solving both simple and complex problems in a timely fashion, while working in conjunction with strategic planning and automated control systems.

III. Tactical Immunized Maneuvering System

The Tactical Immunized Maneuvering System (TIMS) constructs motion-based plans, in the form of maneuver sequences (Fig. 4). These maneuver sequences are composed of one or more autopilot commands, along with the scheduling times for command execution. Each autopilot command consists of a mode identifier and corresponding target. The maneuver selection system contains autopilot mode dependent performance models for predicting the motion-based path of maneuver sequences. These maneuver sequences are constructed from basic piloting maneuvers, which are stored in a maneuver database. Artificial immune algorithms are used to select the appropriate maneuvers from the database, and to augment them as necessary in order to achieve tactical objectives. These objectives can be specified by a strategic planner, autonomous executive, or scientists and ground-based operators. Once these maneuver sequences are generated, they are sent to a specialized autopilot system for execution.

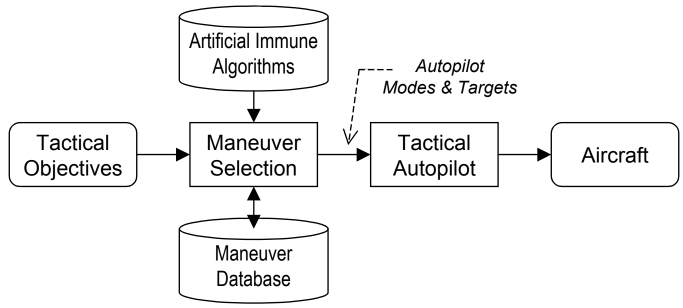

Figure 4. Tactical Immunized Maneuvering System.

A. Immune System Metaphor

A biological immune system can be thought of as a robust adaptive system that is capable of dealing with an enormous variety of disturbances and uncertainties. The AIS combines *a priori* knowledge with the adapting capabilities of a biological immune system to provide a powerful alternative to currently available techniques for pattern recognition, learning and optimization.[9]

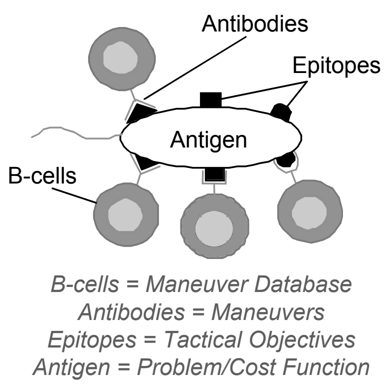

Figure 5. Immune System Metaphor.

In terms of the immune system metaphor (Fig. 5), the infectious agent (or *antigen*) represents the problem and the *antibody* represents the solution. For this application, the problem is expressed in terms of a cost function, and the solution is expressed in terms of maneuvers.

Antigens each have a set of antigenic determinants called *epitopes* that are molecular shapes recognized by the immune system. The antibodies bind with these epitopes to subsequently neutralize them and remove the threat. For this application, the epitopes represent the tactical objectives, which are expressed in terms of cost function parameters.

The antibodies are secreted from *B-cells*, which are produced in the bone marrow. Some of these B-cells survive as memory cells, essentially allowing the solution to the problem to be remembered. For this application, the surviving B-cells represent the successful maneuvers that are stored back into the maneuver database.

B. Immunized Maneuver Selection

During the process of immunized maneuver selection, the search for a solution is modeled after the artificial immune system response (Fig. 6).

Bone Marrow Models:

In bone marrow models, gene libraries are used to create antibodies from the bone marrow. The gene library contains pieces of a solution that has been predetermined using a priori knowledge. Antibodies are produced through a concatenation of genes from the gene library.

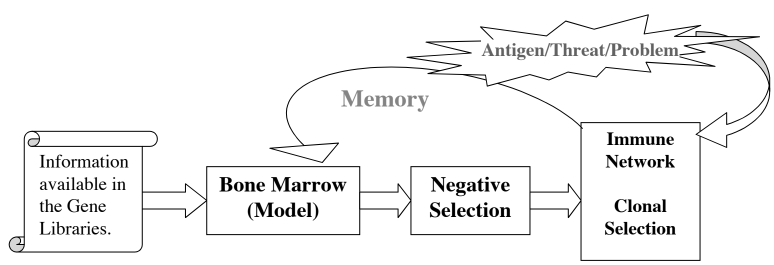

Figure 6. Artificial Immune System Response.

For this application, the binary representations of an autopilot mode, target, and wait time are used to construct basic maneuvers (Fig. 7). These maneuvers can in turn be combined with other maneuvers to form more complex maneuver sequences. The collection of these maneuvers is stored in a maneuver database, and represent a combination of randomly create and/or manually constructed maneuvers, as well as maneuvers that are generated through immunized maneuver selection. In certain situations heuristics can also be used to automatically generate maneuvers, in order to preserve the ability to incorporate rule-based approaches for quickly solving simple problems that are normally encountered.

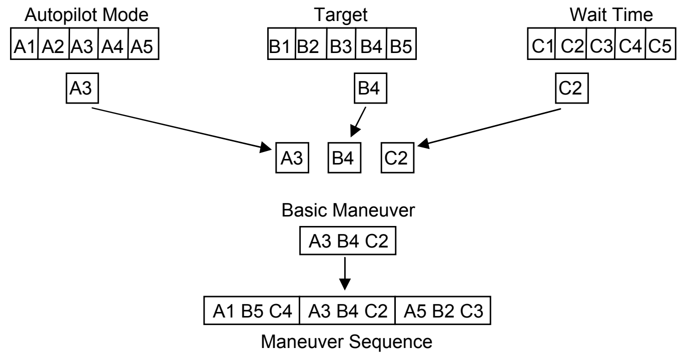

Figure 7. Constructing Maneuver Sequences.

Negative Selection:

Negative selection is based on the principle of self-nonself discrimination in the immune system. This discrimination is achieved in part by *T-cells*, which have receptors on their surface that can detect antigens and activate the necessary B-cells. For this application, this selection process is achieved by comparing the characteristics of the current tactical objectives against the properties of the maneuvers in the maneuver database. The resulting problem-to-solution mapping is stored into a strength matrix. As the connections between tactical objectives and maneuvers grow over time, the likelihood that the necessary maneuvers will be selected from the maneuver database increases. As a result, the time required for finding a solution to a similar problem will be reduced.

Clonal Selection:

The goal of clonal selection in the immune system is to find the most suitable member of a population of antibodies in a very short period of time. The clonal selection algorithm (Fig. 8) uses selection, cloning, and maturation (or hypermutation) to perform the tasks of discovering and maturing good antibodies from the population of available solutions in an orchestrated fashion. This is achieved by allowing antibodies with high performance to have a higher probability of reproduction. Furthermore, as antibodies mature, they are allowed to tuning themselves and thus improve their chances of survival. For this application, the performance of a maneuver is computed by evaluating its predicted motion-based trajectory against the tactical objectives, which are expressed in terms of a cost function. The predicted trajectory is computed using autopilot mode-dependent models.

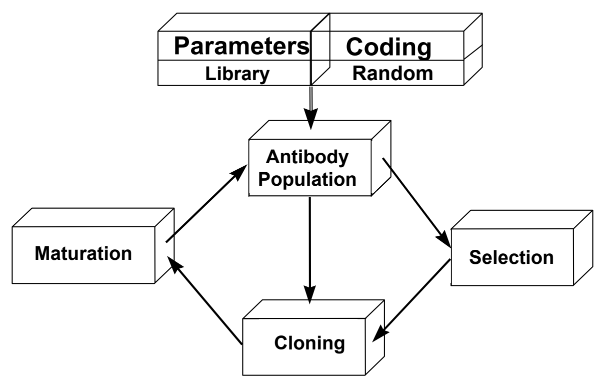

Figure 8. Clonal Selection Algorithm.

Immune Network:

In the immune network theory, antibodies recognize both antigens and other antibodies. Antibodies that recognize other antibodies form a network within the immune system. As the antibody matures, it recognizes the antigen with a higher degree of accuracy. Once the antigen is completely removed, the network between like-antibodies helps in keeping the immune system from extinguishing itself. A stable population is maintained, as a form of memory, so that it will be available for future encounters with similar antigen. For this application, the concept of an immune network is achieved by storing successful maneuvers back into a maneuver database. Furthermore, just as the immune network maintains itself, the maneuver database can also manage itself in order to limit the size of the database. For example, maneuvers that are similar to pre-existing maneuvers may not be stored, while other rarely selected maneuvers may be deleted over time.

C. Tactical Autopilot System

The tactical autopilot system is based upon a neural flight controller and auto-gain scheduling guidance system (Fig. 9), which can be applied to a wide range of vehicle classes.[10] This autopilot has been enhanced with additional modes and an aggressiveness factor for enabling high performance maneuvers. The command interface has also been modified to process mode and target sequences.

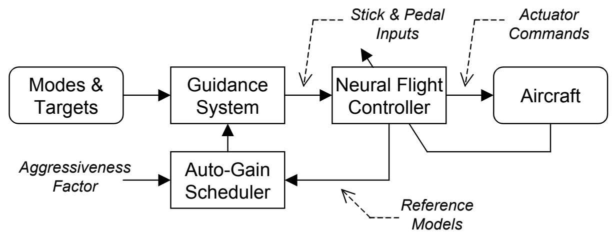

Figure 9. Tactical Autopilot System.

Neural Flight Controller:

The neural flight controller provides consistent handling qualities, across flight conditions and for different aircraft configurations This direct adaptive tracking controller integrates feedback linearization theory with both pre-trained and on-line learning neural networks (Fig. 10).[11] Pre-trained neural networks provide estimates of aerodynamic stability and control characteristics required for model inversion. On-line learning neural networks generate command augmentation signals to compensate for errors in the estimates and from model inversion. Reference models are used to filter inputs in order to shape desired handing qualities.

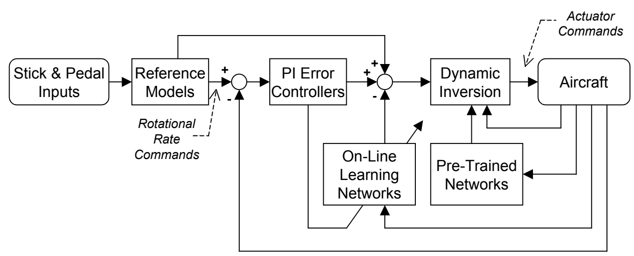

Figure 10. Neural Flight Controller.

Guidance System:

The guidance system takes advantage of the consistent handling qualities in order to achieve deterministic outer-loop performance. Automatic gain-scheduling is performed using frequency separation, based upon an aggressiveness factor and the neural flight controller's specified reference models. The aggressiveness factor is used to limit the percentage of allowable stick and pedal deflections that the guidance system can command. These limits are then propagated throughout the guidance system in the form of computed gains and command limits.

Autopilot commands correspond to control modes, which are based upon a conventional autopilot system. However additional body-axis modes have been added to provide the necessary aerobatic maneuvering capability. Each mode corresponds to control laws, which are built upon each other to form a control hierarchy (Fig. 11). Every autopilot command consists of a mode identifier and corresponding target. Every maneuver sequence is composed of one or more autopilot command, followed by a wait command for scheduling command execution. A semi-colon delimiter separates all commands.

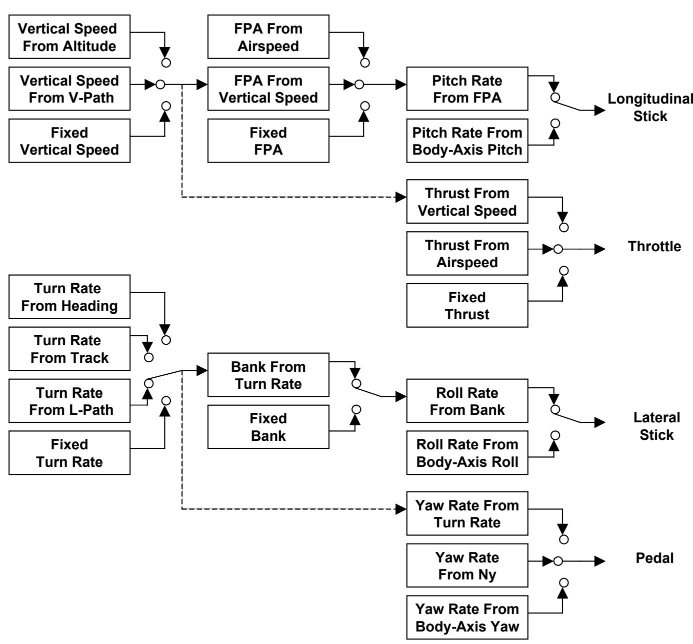

Figure 11. Control Hierarchy.

IV. Implementation and Test Results

Simulation tests were performed to evaluate the potential of using TIMS to enhance exploration missions. Various scenarios were developed to test how the system responds to different tactical objectives under certain target location conditions. This section describes the tests that were performed using a simulated UAV, and the details regarding the corresponding implementation.

A. Simulation Description

Evaluations were conducted using a simulation of the ALTAIR UAV, which is an enlarged version of the Predator B, manufactured by General Atomics – Aeronautical Systems Inc. (GA-ASI). The aircraft was designed with an extended wingspan to perform higher altitude, longer range and extended duration earth science missions for NASA. The simulation model was developed using a Rapid Aircraft Modeler (RAM) to create a three-dimensional representation of the aircraft (Fig. 12a). A "Balance" program was then used to estimate aircraft center-of-gravity (CG) and inertial characteristics. Finally, the aerodynamic stability and control derivatives were computed at different operating conditions using vortex-lattice code (VORVIEW) (Fig. 12b).[12] The six degree-of-freedom flight simulation contains equations of motion that are based on perturbation theory, however most of the non-linear terms are retained in the gravitational and inertial portions of the translational axes. All of the non-linear terms are retained in the inertial portions of the rotational axes. The Earth atmosphere is based on a 1976 standard atmosphere model. The Dryden turbulence model provides turbulence RMS and bandwidth values which are representative those specified in Military Specifications Mil-Spec-8785 D of April 1989.

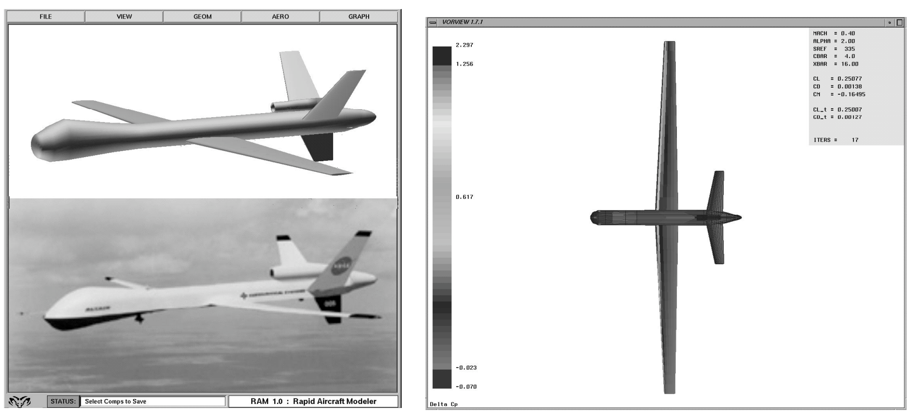

Figure 12.
(a) Three-Dimensional Generated Model (Left);
(b) Mapping of the Surface Pressure Distribution (Right).

The payload sensors were modeled as a downward looking device with a non-pivoting mount. As a result, the vehicle's attitude directly affects the location of the projected data collection footprint. The resolution of the sensor decreases with both the distance of the aircraft from the target, and the distance of the target from the center of the data collection footprint. Therefore the highest data quality can be achieved by flying directly over the target so that the payload sensors are pointing straight down. In order to test the effectiveness of high precision maneuvers, the satisfactory data collection footprint was limited to a 1^{o} field of view, which corresponds to a sea level radius of approximately 350 feet when the vehicle is flying at an altitude of 40000 feet (Fig. 13). However, the sensors were also assumed to be capable of detecting targets with a larger 20^{o} field of view, corresponding to a sea level radius of approximately 14500 feet. Once the sensors detect a target, the updated position information could be used to generate higher precision maneuvers.

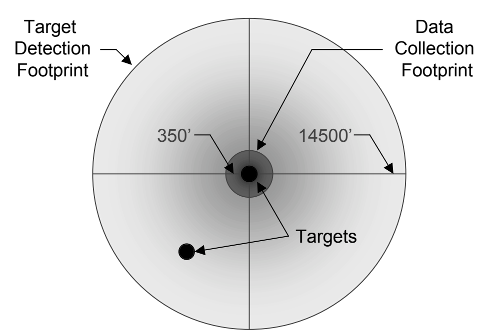

Figure 13. Payload Sensor Footprint.

B. Gene Library Implementation

Each maneuver sequence command consists of at lest one autopilot mode identifier and target, and the associated time to wait after command execution. The genetic representation of a single maneuver is defined by a 32-bit word (Fig. 14). As a result, there can be up to 16 possible autopilot modes. The maximum and minimum target values and wait times, corresponding to each autopilot mode, can be subdivided into up to 16 different regions. The precision factors dictate the resolution of the target values and wait times within each region.

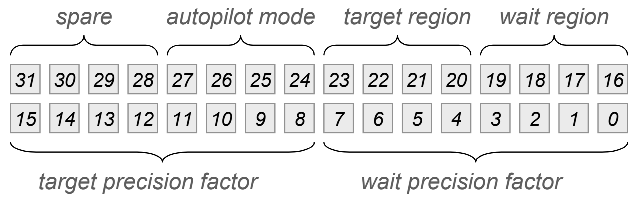

Figure 14. Genetic Representation of a Maneuver.

For this evaluation, the minimum wait times corresponded to 3 times the time constant for each autopilot mode. This allows the modeling routines to assume that the autopilot modes had essentially reached a steady state prior to executing the next command. The maximum wait times corresponded to 17 times the time constant. However, a "do nothing" autopilot mode was added, with a minimum wait time of one second and a maximum wait time of 1 hour, to enable increased attainable wait times between commands.

The number regions for breaking up the range of target values and wait times were limited to 5 in order to reduce the solution space, while still providing adequate precision. An odd number of regions were chosen to ensure that the middle region spans a target value of zero (Fig. 15). The following autopilot modes were used in this evaluation, along with the corresponding maximum and minimum target values:

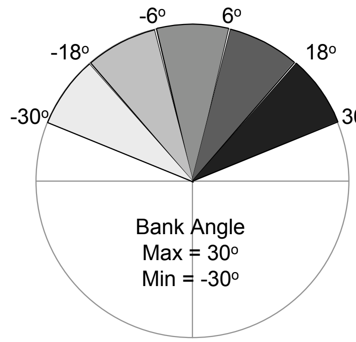

Figure 15. Bank Target Regions.

1) Bank angle (bank) commands the vehicle to maintain a certain bank angle. The maximum and minimum target values (+/- 30 degrees) were selected so that the vehicle can still maintain level turn throughout the maneuver.
2) Delta heading (dhdg) commands the vehicle to a heading that is relative to the heading at the time the command is issued. The maximum and minimum target values (+/- 30 degrees) were selected so that the bank angle limits would not be reached.
3) Flight-path angle (fpa) commands the vehicle to maintain a certain flight-path angle. The maximum and minimum target values (+/- 5 degrees), which correspond to approximately 1300 feet per minute climb and descent rates at the desired flight condition, were selected so that the vehicle can still maintain airspeed.
4) Delta altitude (dalt) commands the vehicle to an altitude that is relative to the altitude at the time the command is issued. The maximum and minimum target values (+/- 500 feet) were selected so that the flight-path angle limits would not be reached.
5) Delta airspeed (dspd) commands the vehicle to maintain a certain indicated airspeed. The maximum and minimum target values (+/- 25 knots) were selected so that the throttle limits would not be reached at the desired flight conditions.
6) No command (none) provides a means for the vehicle to extend the wait time of the previous command.

Autopilot mode dependent performance models were used to predict the flight path of the vehicle during a maneuver sequence. Each mode was modeled in terms of a first order model with the mode-dependent time constants. This allowed the prediction to be performed with very little computation, verses utilizing more accurate, but also computationally intensive methods such as fast-time simulation. While this modeling method provides an estimate of the vehicle's state at the end of each command, it does not account for intermediate states throughout the maneuver such as the proximity to obstacles.

C. Antigen Implementation

The cost function is expressed in terms of weighted parameters, represented by the tactical objectives. These tactical objectives can incorporate both rewards and penalties. For this evaluation, reward (indicated by a negative cost) was applied for servicing a target. The amount of the rewards was computed as a function of the predicted quality of the data, as determined by the distance to the target and the target's position relative to center of the payload sensor footprint (Fig. 13).

Penalties (indicated by a positive cost) were applied for constraint violations, such as not achieving the necessary bank angle and/or heading, at the time of target servicing. For safety purposes, large penalties were applied any time the vehicle was predicted to approach the limits of the allowable flight envelope for the mission. Additional smaller penalties were also applied for the overall duration of the maneuver and the number of maneuvers in the maneuver sequence.

In all of these cases, rewards and penalties were computed as functions of desired verses predicted states, so that the computed performance index could discriminate between slight differences between candidate maneuvers. However, the complex nature of the cost function still results in a high likelihood of getting caught in local minima for an extended period of time. Fortunately, as opposed to traditional hill-climbing techniques, genetic algorithms are designed to move the population away from local minima. Figure 16 shows an example where substantial drops in the performance index are indicative of escaping from local minima situations. As a result, it can be readily apparent when a target becomes serviced.

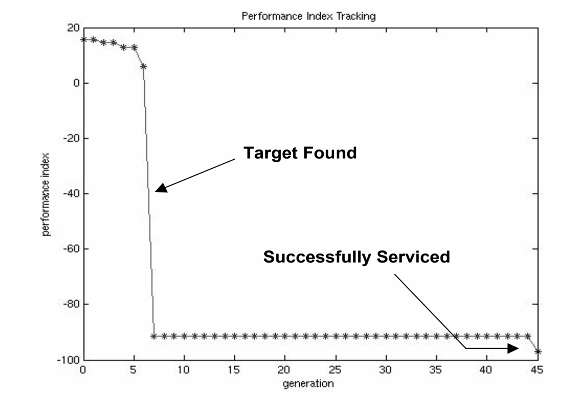

Figure 16. Performance Index Tracking.

For this evaluation, the targets were represented as points verses regions. As a result, the vehicle's states were evaluated at the end of each autopilot command. In cases were obstacles are present, or a vehicle must point to a target for an extended point of time, the vehicle's states would need to be evaluated throughout the maneuver. This could be accomplished by breaking up the predicted motion-based trajectory into smaller time-steps, or by incorporating fast-time simulation capability.

D. Negative Selection Implementation

The strength matrix provides a means of performing negative selection through problem-to-solution mapping. This is achieved by comparing the characteristics of the problem with the properties of the potential solutions. For this evaluation, the characteristics of the problem were determined by categorizing the relative target location and associated servicing constraints into different regions. Similarly, the properties of each solution were determined by categorizing the vehicle's final position and state into the same regions. The strength of the problem-to-solution mapping is then assessed by comparing the number of matching regions, for each maneuver in the maneuver database.

Figure 17 shows an example of the strength matrix that was used for this evaluation. The problem characteristics were organized into 5 separate categories (Δx, Δy, Δh, $\Delta\psi$, ϕ). Each category was divided into 5 different regions.

Problem Characteristics

Solution Properties

Manv. No.	Δx (ft)					Δy (ft)					Δh (ft)					$\Delta\psi$ (deg)					ϕ (deg)				
	Δx--	Δx-	Δx_o	Δx+	Δx++																				
1				X					X				X						X				X		
2			X			X										X					X				
3	X							X																	X
• • •																									
N			X						X					X				X					X		

Figure 17. Strength Matrix.

1) The Δx category is defined to be the distance of the target in front of the aircraft. The Δx regions correspond to the radius of the target detection and target servicing footprints, where (Δx-- < -8700 ft), (-8700 ft $\leq \Delta x$- < -2900 ft), (-2900 ft $\leq \Delta x_o$ < 2900 ft), (2900 ft $\leq \Delta x$+ < 8700 ft), (8700 ft $\leq \Delta x$++).
2) The Δy category is defined to be distance of the target to the right of the aircraft. The Δy regions correspond to the radius of the target detection and target servicing footprints, where (Δy-- < -8700 ft), (-8700 ft $\leq \Delta y$- < -2900 ft), (-2900 ft $\leq \Delta y_o$ < 2900 ft), (2900 ft $\leq \Delta y$+ < 8700 ft), (8700 ft $\leq \Delta y$++).
3) The Δh category is defined to be the relative altitude required for target servicing. The Δh regions correspond to the maneuvering capabilities of the aircraft, where (Δh-- < -20000 ft), (-20000 ft $\leq \Delta h$- < -8000 ft), (-8000 ft $\leq \Delta h_o$ < 8000 ft), (8000 ft $\leq \Delta h$+ < 20000 ft), (20000 ft $\leq \Delta h$++).

4) The $\Delta\psi$ category is defined to be the relative heading required for target servicing, relative to the current heading of the aircraft. The $\Delta\psi$ regions are divided equally, where (-180 deg $\leq \Delta\psi$ -- < -108 deg), (-108 deg $\leq \Delta\psi$- < -36 deg), (-36 deg $\leq \Delta\psi_o$ < 36 deg), (36 deg $\leq \Delta\psi$+ < 108 deg), (108 deg $\leq \Delta\psi$++ < 360 deg).
5) The ϕ category is defined to be the bank angle required for target servicing. The ϕ regions correspond to the bank angle limits of the aircraft, where (-30 deg $\leq \Delta\psi$ -- < -18 deg), (-18 deg $\leq \Delta\psi$- < -6 deg), (-6 deg $\leq \Delta\psi_o$ < 6 deg), (6 deg $\leq \Delta\psi$+ < 12 deg), (12 deg $\leq \Delta\psi$++ < 30 deg).

In the case of multiple targets, the problem characteristics would need to be expanded to include the relative position and constraints of all of the targets. Similarly, the solution properties would also be expanded to include the vehicle's positions and states at the end of each autopilot command throughout a maneuver sequence. As a result, the strength of each maneuver in the maneuver database would be assessed against for the overall constellation of targets.

E. Clonal Selection Implementation

The clonal selection algorithm (Fig. 8) provides a method of finding the most suitable member of a population in a relatively short period of time. For this evaluation, a population of 500 maneuvers was maintained. The number of generations was dependent upon how long it took to settle in on a solution that successfully serviced the desired target. However, in general, an upper limit of 5 computational seconds was applied.

During the selection process, the top 50% of the population was permitted to survive and mature through evolutionary tuning of the precision factors. Another 25% immigrated into the population from the maneuver database, through the negative selection process. The final 25% were reproduced from random members in the population using a combination of standard and uniform genetic mutation, and the random addition or subtraction of maneuvers in the maneuver sequence.

F. Immune Network Implementation

The maneuver database was initialized with a combination of randomly generated and manually constructed maneuvers. The randomly generated maneuvers contained anywhere from 1 to 3 autopilot commands. Then manually constructed maneuvers consisted of a series of S-Turns, bank commands followed by extended wait commands and heading commands, flight-path angle commands followed by extended wait commands and altitude commands, and flight-path angle commands followed by bank commands.

G. Test Results

Simulation tests were performed with the aircraft initialized flying straight and level at an altitude of 40,000 feet, and at a true airspeed of 150 knots. All tests were performed in light turbulence. The results for each test contain both the predicted and actual trajectories for comparison. Since maneuver re-planning was not performed, while an existing maneuver was in the process of being executed, the affect of predictive modeling errors can be readily identified.

Constrained Attitude Test (Maximizing Data Quality):

Mission requirements can often dictate how a vehicle must service a target. These requirements can place constraints on the aircraft's attitude. For example, missions involving target mapping may require the vehicle to maintain a zero bank angle. In some cases, the vehicle's heading may also need to be constrained relative to the sun angle. Another problem is that many times the precise location of the target is not known. As a result, the vehicle may need to make quick corrections to the target once it is detected.

Figure 18a shows a scenario where the vehicle is currently flying on its desired heading in the general direction of a target. However when the target enters the 14500 feet radius of the detection region (Fig. 13), it turns out to be offset by 45 degrees to the right. As a result, the vehicle needs to make a correction of 10000 feet to the right (Δx) before it travels another 10000 feet in the forward direction (Δy), and still maintain a parallel heading with a zero bank angle as it passes over the target. The resulting maneuver sequence commands (bank 23.04; wait 28.09; bank 3.46; wait 10.61; bank –24.54; wait 28.23; bank 0.23; wait 6.05) represent a pseudo S-Turn with a near zero bank command at the end to bring the wings back to level. A total of 28 clonal selection generations were computed in less than 3 seconds, before the successful target servicing performance index was reached. The results contain both predicted and actual trajectories for comparison. Figure 18b shows the corresponding predicted and actual bank angles throughout the maneuver. The predicted bank angle was fairly accurate, although the actual bank angle tended to slightly lag the prediction by about a second. This effect, along with turn coordination inaccuracies in the simulated aircraft, resulted in a ground track prediction that was off by approximately 1500 feet by the end of the maneuver. However, the vehicle did successfully correct for the lateral offset and was in position to service the target.

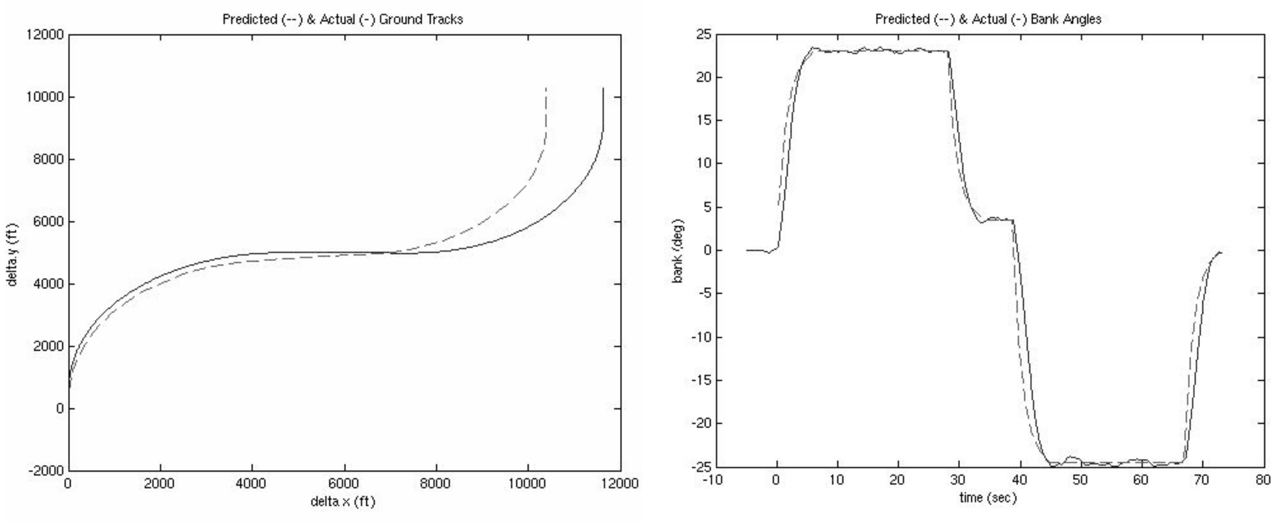

Figure 18.
(a) Predicted (--) & Actual (-) Ground Tracks (Left);
(b) Predicted (--) & Actual (-) Bank Angles (Right).

Target Pointing Test (Minimizing Time):

Another unique aspect of exploration missions is that the primary focus is in regards to the projection of the payload sensor, rather the position of the vehicle. As a result, a vehicle may need to be able to point a sensor in the direction of a target. This becomes especially crucial in a target rich environment, or during situations where there may be insufficient time to fly to every target.

Figure 19 shows a scenario where a target enters the very edge of the detection region. In this case, the target is 90 degrees to the right, which means that the vehicle needs to make a correction of 14500 feet to the right (Δx) and 0 feet in the forward direction (Δy). The only way the vehicle can fly to the target is to turn around and fly back towards the target (bank 27.07; wait 28.33; dhdg 21.24; wait 45.24). A total of 45 generations were computed in less than 5 seconds, before the successful target servicing performance index was reached. Once again the predicted trajectory passes over the target, and the actual trajectory comes to within 1500 feet of the target. The entire maneuver took over a minute (73.57 seconds).

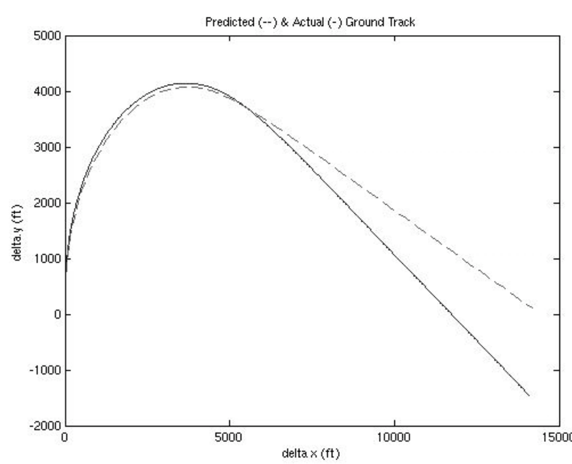

Figure 19. Predicted (--) & Actual (-) Ground Tracks Back to Target.

Figure 20a shows the same scenario, however the time penalty in the tactical objective was increased. The resulting maneuver (bank 25.80; wait 8.17; bank –19.70; wait 6.82) took just under 15 seconds. In this case, instead of flying over the target, the vehicle maneuvers in such a way that the resulting sensor footprint is projected over the target. However, since this involves a projection of the vehicle's predicted attitude, the predicted verses actual footprint errors are significantly larger (Fig. 20b). Although the predicted and actual vehicle positions and attitudes appear close to each other, the two sensor footprints are almost 4000 feet apart by the end of the maneuver.

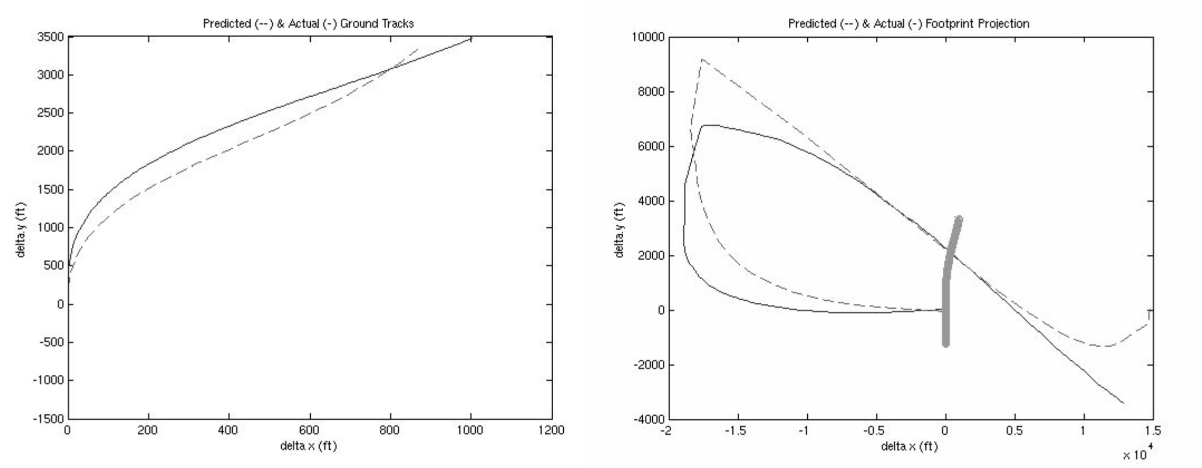

Figure 20.
(a) Predicted (--) & Actual (-) Ground Tracks (Left);
(b) Predicted (--) & Actual (-) Sensor Footprints (Right).

Learning Test (Reducing Computations Over Time):

In order for intelligent systems to assume increasingly larger roles in semi-autonomous and autonomous missions, they must incorporate the experience, reasoning and learning abilities of a pilot. One method TIMS uses to incorporate a pilot's experience is through the manual construction of typical piloting maneuvers, which are then stored in the maneuver database. Additionally, Figures 18 through 20 show the results of the predictive model-based reasoning aspects of the clonal selection process. However, perhaps one of the most important human characteristics is the ability to learn through experience. This ability is incorporated into TIMS through a combination of storing successful maneuvers into the maneuver database, as part of the immune network model, and by characterizing the problem-to-solution relationship through the strength matrix, as part of the negative selection process.

In order to demonstrate the learning process, the pre-constructed maneuvers were removed from the maneuver database. Then a series of random targets were placed within the detection footprint. The system was then responsible for generating maneuvers in order to service those targets by flying over them with wings-level. This method represents a form of off-line learning that can be used to pre-train the system for a given class of mission. Figure 21a shows a plot of the number of generations that were required before a satisfactory solution was found.

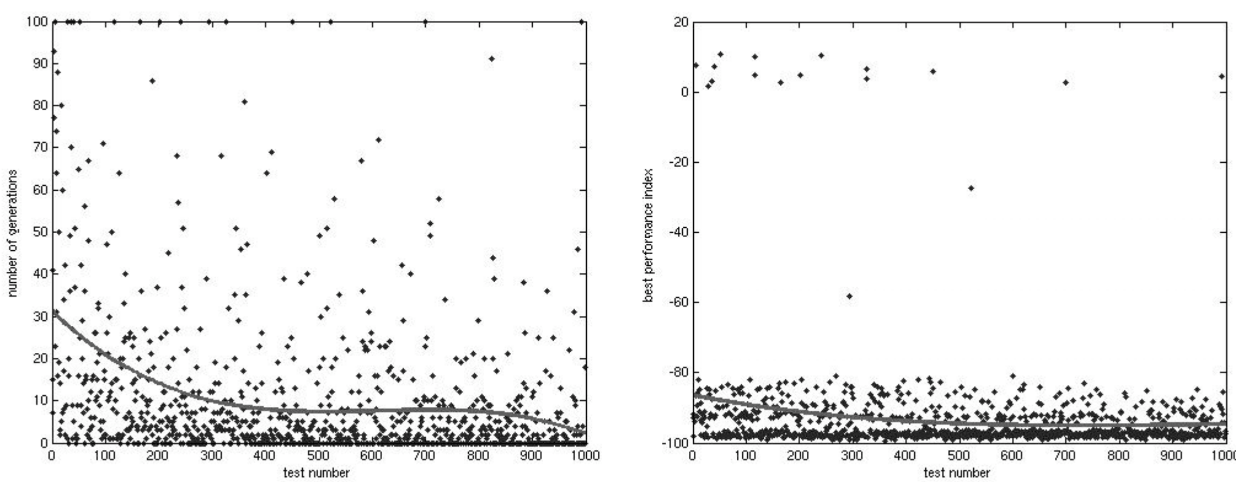

Figure 21.
(a) Number of Generations Over Time (Left);
(b) Best Performance Index Over Time (Right).

Figure 21b shows a plot of the performance indices of the satisfactory solutions. While both the number of generations and the performance indices varied throughout the test, as a result of the random level of problem difficulty, the overall number of generations and performance indices decreased as the system learned over time.

V. Conclusion

By applying intelligent methods of onboard automation, pilots and ground-based operators can defer many of the responsibilities of performing and supervising tasks to focus on managing goals and objectives. This becomes especially essential for exploration UAVs that will typically be operating in remote or planetary environments. Since the goal of exploration is to collect data, one of the responsibilities of the onboard intelligence will be to react to the data with little or no human direction.

There are many of different methods of reacting to data. Strategic planners can provide excellent tracking, and are especially applicable to missions were the target locations are either known or can be determined well in advance by the sensors. The role of tactical maneuvering is oriented more towards providing a means of replacing the missing piloting skills. Unlike flight management systems, pilots are not completely dependent upon precise navigational information. They have the ability to assess a situation, with only partial information, in order to select the best near-term course of action.

The immunized maneuver selection approach has demonstrated the potential of solving difficult problems within a matter of seconds. Even more importantly, the system has the ability to learn by becoming familiar with the best ways to operate for a particular class of mission. However, as the planning process transitions from near-term to longer-term focuses, errors in model prediction can become apparent. Fortunately, one method of handling this occurrence is by taking the pilot oriented approach of constantly monitoring a situation in order to revise the near-term plan if necessary. As a result, future research focuses could potentially include the following areas:

1) Tactical re-planning in order to mature (or replace) existing maneuvers by incorporating existing autopilot modes and maneuver sequence commands into the decision making process. This could be addressed by mutating the existing plan, or simply maturing the target values and wait times of the existing command.
2) Incorporating alternate fast-time simulation approaches to improve prediction forecasting and to enable transient behavior modeling. The primary tradeoff is that additional computational steps results in longer solution generation times. However it may be possible to offset some of these factors by incorporating a type of regulated population control and target servicing criteria. Further development is also necessary in the area of maneuver database management.
3) Expansion of the target definitions to include multi-dimensional target regions and constraints, and to incorporate the ability to track moving targets and to assess multiple target constellations.

References

[1]Wegener, Steven S., Schoenung, S., Totah, J., Sullivan, D., Frank, J., Enomoto, F., Frost, C., and Theodore, C., "UAV Autonomous Operations for Airborne Science Missions," AIAA-2004-6416

[2]KrishnaKumar, K., Kaneshige, J., and Satyadas, A., "Challenging Aerospace Problems for Intelligent Systems," *Proceedings of the von Karman Lecture series on Intelligent Systems for Aeronautics*, Belgium, May 2002.

[3]Chen, Ted L., and Pritchett, Amy R., "On-The-Fly Procedure Development for Flight Re-Planning Following System Failures," AIAA-2000-0300.

[4]Rathbun, D., and Capozzi, B., "An Evolution Based Path Planning Algorithm for Autonomous Motion of a UAV through Uncertain Environments," *Proceedings of the 21st Digital Avionics Systems Conference (8D2),* Irvine, CA, October 2002.

[5]Judd, Kevin B., and McLain, Timothy W., "Spline Based Path Planning For Unmanned Air Vehicles," AIAA-2001-4238.

[6]Gavrilets, V., Martinos, I., Mettler, B., and Feron, E., "Flight Tests and Simulation Results for an Autonomous Aerobatic Helicopter," *Proceedings of the 21st Digital Avionics Systems Conference (8C3),* Irvine, CA, October 2002.

[7]Frazzoli, E., "Maneuver-Based Motion Planning and Coordination for Multiple UAV's," *Proceedings of the 21st Digital Avionics Systems Conference (8D3),* Irvine, CA, October 2002.

[8]Kaneshige, J., KrishnaKumar, K., and Shung, F., "Tactical Maneuvering Using Immunized Sequence Selection," AIAA-2003-6640.

[9]Krishnakumar, K., "Artificial Immune System Approaches for Aerospace Applications," AIAA-2003-0457.

[10]Kaneshige, J., Bull, J., and Totah, Joseph J., "Generic Neural Flight Control and Autopilot System," AIAA-2000-4281.

[11]Rysdyk, Rolf T., and Calise, Anthony J., "Fault Tolerant Flight Control via Adaptive Neural Network Augmentation," AIAA-98-4483.

[12]Totah, Joseph J., Kinney, David J., Kaneshige, John T., and Agabon, Shane, "An Integrated Vehicle Modeling Environment," AIAA-99-4106.